arXiv:1704.02407v1 [math.CO] 8 Apr 2017

MORE ON ADDITIVE TRIPLES OF BIJECTIONS

SEAN EBERHARD

Abstract. We study additive properties of the set S of bijections (or permutations) $\{1,\dots,n\} \to G$, thought of as a subset of G^n, where G is an arbitrary abelian group of order n. Our main result is an asymptotic for the number of solutions to $\pi_1 + \pi_2 + \pi_3 = f$ with $\pi_1, \pi_2, \pi_3 \in S$, where $f : \{1,\dots,n\} \to G$ is an arbitary function satisfying $\sum_{i=1}^n f(i) = \sum G$. This extends recent work of Manners, Mrazović, and the author [EMM15]. Using the same method we also prove a less interesting asymptotic for solutions to $\pi_1 + \pi_2 + \pi_3 + \pi_4 = f$, and we also show that the distribution $\pi_1 + \pi_2$ is close to flat in L^2.

As in [EMM15], our method is based on Fourier analysis, and we prove our results by carefully carving up $\widehat{G}^n$ and bounding various character sums. This is most complicated when G has even order, say when $G = \mathbf{F}_2^d$.

At the end of the paper we explain two applications, one coming from the Latin squares literature (counting transversals in Latin hypercubes) and one from cryptography (PRP-to-PRF conversion).

1. Introduction

Let G be an abelian group of order n, and let S be the set of bijections $\pi : \{1,\dots,n\} \to G$. We are interested in additive properties of S as a subset of G^n. For example, we are interested in counting solutions to

$$\pi_1 + \pi_2 = \pi_3 \qquad (\pi_1, \pi_2, \pi_3 \in S).$$

The solutions to this equation are called *additive triples* (or *Schur triples*). The following estimate was the main theorem of [EMM15].

Theorem 1.1 (Main theorem of [EMM15])**.** *Assume n is odd. Then the the number of additive triples in S is*

$$(e^{-1/2} + o(1))n!^3/n^{n-1}.$$

The odd-order hypothesis here is irritating. In [EMM15] our main interest was the cyclic group $G = \mathbf{Z}/n\mathbf{Z}$, and in that case the restriction to odd n is natural, since otherwise there aren't any solutions to $\pi_1 + \pi_2 = \pi_3$, but in general it is a shortcoming. It is a little sad not to have an asymptotic in the case $G = \mathbf{F}_2^d$, for example. Here, among other things, we make amends by proving the following theorem.

Theorem 1.2. *Assume $\sum G = 0$. Then the number of additive triples in S is*

$$(e^{-1/2} + o(1))n!^3/n^{n-1}.$$

Here and throughout the paper we write simply $\sum G$ for $\sum_{x \in G} x$, which is a characteristic element of G. It is easy to see that $\sum G$ is equal to the unique element of order 2 in G if there is one, and otherwise zero. Thus $\sum G \neq 0$ if and only if G has the form $\mathbf{Z}/m\mathbf{Z} \times H$ with m even and $|H|$ odd. In these groups

there are no additive triples in S (since $\pi_1 + \pi_2 = \pi_3$ implies $2\sum G = \sum G$), and Theorem 1.2 covers all other groups.

Generalizing the problem, suppose we count solutions to

$$\pi_1 + \pi_2 + \pi_3 = f,$$

where $f : \{1, \dots, n\} \to G$ is a fixed but arbitrary function. (The previous problem is the case $f = 0$, since π is a bijection if and only if $-\pi$ is a bijection.) Call such solutions *additive f-triples*. The natural hypothesis about f and G is that $\sum_{i=1}^n f(i) = \sum G$. Our main theorem is the following generalization of Theorem 1.2.

Theorem 1.3. *Let $f : \{1, \dots, n\} \to G$ be a function satisfying $\sum_{i=1}^n f(i) = \sum G$. Then the number of additive f-triples in S is*

$$(\mathfrak{S}(f) + o(1))n!^3/n^{n-1},$$

where

$$\mathfrak{S}(f) = \exp\left(-\frac{1}{2n^2}\sum_{x \in G} |f^{-1}(x)|^2\right).$$

By analogy with analytic number theory, we might call $\mathfrak{S}(f)$ the *singular series* associated with f. The expression inside the brackets is closely related to the so-called *collision entropy* (or *Rényi entropy*) of f, defined by

$$H_2(f) = -\log \sum_{x \in G} \mathbf{P}(f = x)^2.$$

Explicitly,

$$\mathfrak{S}(f) = \exp\left(-\frac{1}{2}e^{-H_2(f)}\right).$$

The value of $\mathfrak{S}(f)$ can be as small as $e^{-1/2}$ (when f is constant), but generically $\mathfrak{S}(f) = 1 + o(1)$. Thus we can describe the distribution of $\pi_1 + \pi_2 + \pi_3$ as pretty close to uniform on the coset $\sum_{i=1}^n f(i) = \sum G$, but with a slight aversion to functions with high collision entropy.

If we have four or more π summands, then distribution becomes asymptotically flat. In fact this follows already from Theorem 1.3, but by directly applying our method we will prove the following quantitative version of this assertion.

Theorem 1.4. *Let $d \geq 4$ be an integer, and let $f : \{1, \dots, n\} \to G$ a function satisfying $\sum_{i=1}^n f(i) = d\sum G$. Then the number of solutions to*

$$\pi_1 + \cdots + \pi_d = f \qquad (\pi_1, \dots, \pi_d \in S)$$

is

$$\left(1 + O_d(n^{3-d})\right) n!^d/n^{n-1}.$$

Finally, suppose we have just two bijections π_1 and π_2, and we consider $\pi_1 + \pi_2$. In the case of just two bijections, the distribution of is very far from flat in any L^∞ sense. Indeed, it follows from Theorem 1.2 and a symmetry argument that (provided $\sum G = 0$) the number of solutions to $\pi_1 + \pi_2 = \pi$, where π is any fixed bijection, is $(e^{-1/2} + o(1))n!^2/n^{n-1}$, while clearly the number of solutions to $\pi_1 + \pi_2 = 0$ is $n!$. However, we can at least prove that $\pi_1 + \pi_2$ is close to uniform (on the coset $\sum_{i=1}^n f(i) = 0$) in L^2. The following theorem is a slight generalization.

Theorem 1.5. *Let $m < n$ be an integer, and let*

$$\pi_1, \pi_2 : \{1, \dots, m\} \to G$$

be uniformly random injections. Then the L^2 distance between the distribution of $\pi_1 + \pi_2$ and the uniform distribution is $O(m/n^{3/2})$, where the L^2 norm is taken with respect to the uniform distribution. In other words, if S_m is the set of injections then

$$\left\| \left(\frac{n^m (n-m)!}{n!} \right)^2 1_{S_m} * 1_{S_m} - 1 \right\|_2 \leq O\left(m/n^{3/2} \right).$$

Thus in particular the total variation distance between the distribution of $\pi_1 + \pi_2$ and the uniform distribution is $O(m/n^{3/2})$.

At the end of the paper we give two applications. The first is to counting transversals in Latin hypercubes. Just as counting solutions to $\pi_1 + \pi_2 = \pi_3$ is equivalent to counting transversals in a certain Latin square, so counting solutions to $\pi_1 + \cdots + \pi_d = \pi_{d+1}$ is equivalent to counting transversals in a certain Latin hypercube. Thus our estimates answer some (modest) questions in the Latin squares literature.

The second application comes from cryptography, and has to do with conversion of pseudorandom permutations (PRPs) to pseudorandom functions (PRFs). A common construction is to take two pseudorandom permutations $\{0,1\}^d \to \{0,1\}^d$ and to use their bitwise xor. The security of this construction is closely related to Theorem 1.5 with $G = \mathbf{F}_2^d$. In the cryptographers' language, Theorem 1.5 implies that the advantage to an adversary with access to at most m queries is at most $O(m/2^{3d/2})$, as long as $m < 2^d$.

2. Overview of the paper

As usual in additive combinatorics, our main tool for counting additive triples is the Fourier transform. We fix our conventions by giving G the uniform measure and $\widehat{G}$ the counting measure. Thus the Fourier transform $\widehat{1_S}$ of 1_S is defined by

$$\widehat{1_S}(\chi) = \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' \chi_1(x_1) \cdots \chi_n(x_n) \qquad (\chi = (\chi_1, \dots, \chi_n) \in \widehat{G}^n),$$

where the sum is primed to indicate that only distinct $x_1, \dots, x_n \in G$ are to be considered. The number of additive f-triples in S can then be expressed as

$$n^{2n} 1_S * 1_S * 1_S(f) = n^{2n} \sum_{\chi \in \widehat{G}^n} \widehat{1_S}(\chi)^3 \chi(f).$$

In [EMM15] we estimated this sum (in the case $f = 0$, in which case $\chi(f) = 1$ identically) by distinguishing various regions of $\widehat{G}^n$ depending on the *entropy* $H(\chi)$ of $\chi \in \widehat{G}^n$, which we defined as follows. Suppose χ has a_1 coordinates equal to ψ_1, a_2 coordinates equal to ψ_2, and so on, where $\psi_1, \dots, \psi_k$ are distinct and $\sum_{i=1}^k a_i = n$. Since $\widehat{1_S}(\chi)$ is a totally symmetric function of $\chi_1, \dots, \chi_n$, it makes sense to denote $\widehat{1_S}(\chi)$ with the simple shorthand

$$\widehat{1_S}(\psi_1^{a_1}, \dots, \psi_k^{a_k}).$$

The number of χ so represented is

$$\binom{n}{a_1, \dots, a_k}.$$

We define $H(\chi)$ by

$$H(\chi) = \frac{1}{n} \log \binom{n}{a_1, \dots, a_k}.$$

The language is motivated by the fact that $H(\chi)$ is approximately the (Shannon) entropy of the random variable which takes value i with probability a_i/n:

$$H(\chi) \approx \sum_{i=1}^{k} \frac{a_i}{n} \log \frac{n}{a_i}.$$

(See Lemma 5.1 for a precise version of this approximation.) Our basic division of $\widehat{G}^n$ is then defined by $H \leq \varepsilon$ and $H > \varepsilon$, where ε is a small positive parameter:

$$\sum_{\chi \in \widehat{G}^n} \widehat{1_S}(\chi)^3 \chi(f) = \sum_{H(\chi) \leq \varepsilon} \widehat{1_S}(\chi)^3 \chi(f) + \sum_{H(\chi) > \varepsilon} \widehat{1_S}(\chi)^3 \chi(f).$$

The sum over low-entropy characters can in turn be related to a sum over *sparse* characters, where we call χ m-sparse if exactly m of its coordinates are nonzero. First, by a straightforward calculation, if $H(\chi) \leq \varepsilon$ then some coordinate of χ is repeated at least $(1 - 2\varepsilon)n$ times. Second, assuming $\sum_{i=1}^{n} f(i) = \sum G$, the expression $\widehat{1_S}(\chi)^3 \chi(f)$ is invariant under global shifts of χ: if $\chi_i' = \chi_i + \psi$ for each i then

$$\begin{aligned} \widehat{1_S}(\chi') &= \frac{1}{n^n} \sum_{x_1, \dots, x_n}{}' \chi_1(x_1) \cdots \chi_n(x_n) \psi(x_1 + \cdots + x_n) \\ &= \widehat{1_S}(\chi) \psi\left(\textstyle\sum G\right), \end{aligned} \tag{2.1}$$

so, since $2\sum G = 0$,

$$\widehat{1_S}(\chi')^3 \chi'(f) = \widehat{1_S}(\chi)^3 \chi(f) \psi(3\textstyle\sum G + \sum f) = \widehat{1_S}(\chi)^3 \chi(f).$$

Thus, by always shifting so that the majority coordinate is 0, we can decompose the sum over low-entropy characters as

$$\sum_{H(\chi) \leq \varepsilon} \widehat{1_S}(\chi)^3 \chi(f) = n \sum_{m=0}^{2\varepsilon n} \sum_{\substack{m\text{-sparse } \chi \\ H(\chi) \leq \varepsilon}} \widehat{1_S}(\chi)^3 \chi(f).$$

Thus we have

$$\begin{aligned} \sum_{\chi \in \widehat{G}^n} \widehat{1_S}(\chi)^3 \chi(f) &= n \sum_{m=0}^{2\varepsilon n} \sum_{\substack{m\text{-sparse } \chi \\ H(\chi) \leq \varepsilon}} \widehat{1_S}(\chi)^3 \chi(f) + \sum_{H(\chi) > \varepsilon} \widehat{1_S}(\chi)^3 \chi(f) \\ &= n \sum_{m=0}^{M} \sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 \chi(f) + O\left(n \sum_{m=M+1}^{2\varepsilon n} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \right) \\ &\quad + O\left(\sum_{H(\chi) > \varepsilon} |\widehat{1_S}(\chi)|^3 \right). \end{aligned}$$

Here M is a parameter satisfying $1 \ll M \ll \varepsilon n$. By analogy with the circle method from analytic number theory, we refer to the terms of the first sum as *major arcs*, and everything else as *minor arcs*.

In [EMM15] we proved a number of bounds for $\widehat{1_S}(\chi)$, or for sums of $\widehat{1_S}(\chi)$, which allowed us to prove satisfactory estimates for these terms in the case of $G = \mathbf{Z}/n\mathbf{Z}$ with n odd. We made clear our position that there are only notational modifications when G is an arbitrary abelian group of odd order, but there were one or two places where we really did need the full strength of the odd order hypothesis. We state the main results again here, with clearly defined hypotheses. (Although we are stating these results with more general hypotheses than in [EMM15], we need not give new proofs: they are the same.)

Theorem 2.1 (Major arcs estimate, Theorem 3.1 from [EMM15])**.** *Let G be an abelian group of order n. If m is even then*

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = \frac{(-1)^{m/2}}{2^{m/2}(m/2)!}\left(\frac{n!}{n^n}\right)^3 + O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right),$$

while if m is odd then

$$\sum_{m\text{-sparse } \chi} \widehat{1_S}(\chi)^3 = O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right).$$

Theorem 2.2 (L^∞ bound for sparse characters, Proposition 6.1 from [EMM15])**.** *Let G be an abelian group of order n, where n is odd. If $m \leq n/3$ then*

$$\max_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)| \leq e^{O(m^{3/2}/n^{1/2}+m^{1/2})}2^{-m/2}\binom{n}{m}^{-1/2}\frac{n!}{n^n}.$$

Theorem 2.3 (Sparseval, Theorem 5.1 from [EMM15])**.** *Let G be an abelian group of order n. If $m \leq n/2$ then*

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^2 \leq O(m^{1/4})e^{O(m^{3/2}/n^{1/2})}\binom{n}{m}^{1/2}\left(\frac{n!}{n^n}\right)^2.$$

Theorem 2.4 (Square-root cancellation, Theorem 4.1 from [EMM15])**.** *Let G be an abelian group of order n. If $\chi = (\psi_1^{a_1}, \dots, \psi_k^{a_k})$, then*

$$|\widehat{1_S}(\chi)| \leq \binom{n+k-1}{k-1}^{1/2}\binom{n}{a_1,\dots,a_k}^{-1/2}\frac{n!}{n^n}.$$

By combining Theorems 2.1 to 2.4 appropriately we proved, for odd-order G and sufficiently small but constant $\varepsilon > 0$,

$$\sum_{m=0}^{2\varepsilon n}\sum_{\substack{m\text{-sparse } \chi \\ H(\chi)\leq\varepsilon}} \widehat{1_S}(\chi)^3 = (e^{-1/2}+o(1))\left(\frac{n!}{n^n}\right)^3$$

and

$$\sum_{H(\chi)>\varepsilon} |\widehat{1_S}(\chi)|^3 = O_\varepsilon\left(\exp(-c\varepsilon n)\left(\frac{n!}{n^n}\right)^3\right), \tag{2.2}$$

and this proves Theorem 1.1.

In this paper we follow the same broad strategy, but the following amendments are necessary.

1. We need a "twisted" version of the major arcs estimate (Theorem 2.1) which includes the factor $\chi(f)$. This is a straightforward modification, but it requires rehashing the proof of that estimate. This is where the singular series $\mathfrak{S}(f)$ series comes from. See Section 3.
2. We need a replacement for the L^∞ bound (Theorem 2.2) which does not require the odd-order hypothesis. Our replacement will however only be effective for m up to roughly $n/(\log n)^2$. This is the most technical part of the paper. See Section 4.
3. Because our replacement for Theorem 2.2 is only effective for m up to $n/(\log n)^2$ we need to revisit the proof of (2.2) in order to cover a wider entropy range. See Section 5.

Having done all this we will then be in a position to prove our main theorems in Section 6. The two applications mentioned in the introduction are covered in Sections 7 and 8 respectively.

3. Major arcs

In this section we prove the following theorem, which generalizes and replaces Theorem 2.1.

Theorem 3.1 (Major arcs estimate)**.** *Let G be a group of order n. If m is even then*

$$\sum_{m\text{-sparse }\chi} \widehat{1_S}(\chi)^3\chi(f) = \frac{1}{(m/2)!}\left(-\frac{1}{2n^2}\sum_{x\in G}|f^{-1}(x)|^2\right)^{m/2}\left(\frac{n!}{n^n}\right)^3 + O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right),$$

while if m is odd then

$$\sum_{m\text{-sparse }\chi} |\widehat{1_S}(\chi)|^3 = O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right).$$

To prove this we recall some formulae from [EMM15, Section 3]. Given an m-sparse character

$$\chi = (\chi_1, \dots, \chi_m, 0^{n-m}) \in \widehat{G}^n,$$

we say that a partition $\mathcal{P}$ of $\{1, \dots, m\}$ *kills* $(\chi_1, \dots, \chi_m)$ if $\sum_{i\in P}\chi_i = 0$ for each part $P \in \mathcal{P}$. Since we are assuming that each χ_i is nonzero, a killing partition of $(\chi_1, \dots, \chi_m)$ can have at most $\lfloor m/2 \rfloor$ parts. Moreover we have the following two calculations (see (3.3) and (3.4) from [EMM15]):

- Suppose every killing partition of $(\chi_1, \dots, \chi_m)$ has at most k parts. Then
$$|\widehat{1_S}(\chi_1, \dots, \chi_m, 0^{n-m})| = O_m\left(\frac{1}{n^{m-k}}\frac{n!}{n^n}\right). \tag{3.1}$$
- Suppose that m is even and that $(\chi_1, \dots, \chi_m)$ is killed by a unique partition with $m/2$ parts. Then
$$\widehat{1_S}(\chi_1, \dots, \chi_m, 0^{n-m}) = \frac{(-1)^{m/2}}{n^{m/2}}\frac{n!}{n^n} + O_m\left(\frac{1}{n^{m/2+1}}\frac{n!}{n^n}\right). \tag{3.2}$$

Our first claim is that we may ignore all χ except those whose nonzero coordinates are killed by a unique partition with $m/2$ parts, i.e., exactly those χ covered by (3.2). Note this already proves the odd case of Theorem 3.1.

Lemma 3.2. *Let $\mathfrak{M}_m \subset \widehat{G}^n$ be the set of all m-sparse χ whose nonzero coordinates are killed by a unique partition with $m/2$ parts. Then*

$$\sum_{m\text{-sparse } \chi \notin \mathfrak{M}_m} |\widehat{1_S}(\chi)|^3 = O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right).$$

Proof. By permutation-invariance we have

$$\sum_{m\text{-sparse } \chi \notin \mathfrak{M}_m} |\widehat{1_S}(\chi)|^3 = \binom{n}{m} \sum_{\substack{\chi_1,\dots,\chi_m \neq 0 \\ (\chi_1,\dots,\chi_m,0^{n-m}) \notin \mathfrak{M}_m}} |\widehat{1_S}(\chi_1,\dots,\chi_m,0,\dots,0)|^3.$$

We now divide the sum up acoording to the partition $\mathcal{P}$ of maximal size which kills $(\chi_1,\dots,\chi_m)$ (if there are ties, just choose one). Since there are only $O_m(1)$ such partitions, we may consider each partition individually. Fix such a partition $\mathcal{P}$ of $\{1,\dots,m\}$, and suppose $\mathcal{P}$ has k parts.

First consider the case $k < m/2$. The number of $(\chi_1,\dots,\chi_m)$ killed by $\mathcal{P}$ is $O_m(n^{m-k})$, so by (3.1) the contribution to the sum from $\mathcal{P}$ is

$$\binom{n}{m} O_m(n^{m-k}) O_m\left(\frac{1}{n^{m-k}}\frac{n!}{n^n}\right)^3 = O_m\left(\frac{1}{n^{m-2k}}\left(\frac{n!}{n^n}\right)^3\right) = O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right).$$

Now consider the case $k = m/2$. By definition of $\mathfrak{M}_m$, every $\chi \notin \mathfrak{M}_m$ killed by $\mathcal{P}$ is also killed by some other partition with $m/2$ parts. Since the number of partitions is $O_m(1)$, the number of such χ is $O_m(n^{m/2-1})$, so again we get a satisfactory bound for the contribution from $\mathcal{P}$:

$$\binom{n}{m} O_m(n^{m/2-1}) O_m\left(\frac{1}{n^{m/2}}\frac{n!}{n^n}\right)^3 = O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right). \qquad \square$$

It remains to estimate the sum

$$\sum_{\chi \in \mathfrak{M}_m} \widehat{1_S}(\chi)^3 \chi(f).$$

Applying (3.2), we have

$$\begin{aligned}\sum_{\chi \in \mathfrak{M}_m} \widehat{1_S}(\chi)^3 \chi(f) &= \sum_{\chi \in \mathfrak{M}_m} \left(\frac{(-1)^{m/2}}{n^{m/2}} + O_m\left(\frac{1}{n^{m/2+1}}\right)\right)^3 \left(\frac{n!}{n^n}\right)^3 \chi(f) \\ &= \frac{(-1)^{m/2}}{n^{3m/2}} \left(\sum_{\chi \in \mathfrak{M}_m} \chi(f)\right) \left(\frac{n!}{n^n}\right)^3 + O_m\left(\frac{|\mathfrak{M}_m|}{n^{3m/2+1}}\left(\frac{n!}{n^n}\right)^3\right).\end{aligned}$$

Since $\mathfrak{M}_m$ has size $\binom{n}{m} O_m(n^{m/2}) = O_m(n^{3m/2})$, this is

$$\frac{(-1)^{m/2}}{n^{3m/2}} \left(\sum_{\chi \in \mathfrak{M}_m} \chi(f)\right) \left(\frac{n!}{n^n}\right)^3 + O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right), \tag{3.3}$$

so we may now concentrate on the sum

$$\sum_{\chi \in \mathfrak{M}_m} \chi(f).$$

We can write

$$\sum_{\chi \in \mathfrak{M}_m} \chi(f) = \sum_{|N|=m} \sum_{\mathcal{P}} \sum_{\chi} \chi(f),$$

where the first sum runs over all subsets $N \subset \{1, \dots, n\}$ of size m, the second sums runs over all partitions $\mathcal{P}$ of N into $m/2$ pairs, and the third sum runs over all χ with nonzero coordinates exactly on N and killed by the partition $\mathcal{P}$ and no other. For illustration consider the case $N = \{1, \dots, m\}$ and

$$\mathcal{P} = \{\{1,2\}, \dots, \{m-1, m\}\}.$$

Then the inside sum runs over all choices of

$$\chi = (\chi_1, \dots, \chi_m, 0^{n-m})$$

such that $\chi_1, \dots, \chi_m$ are nonzero, $\chi_{2i-1} + \chi_{2i} = 0$ for $i = 1, \dots, m/2$, and moreover such that $\chi_k + \chi_l \neq 0$ unless $\{k, l\} = \{2i-1, 2i\}$ for some i. However, at the cost of an error of size only $O_m(n^{m/2-1})$ we can just consider the sum over all $\chi_1, \dots, \chi_m$ such that $\chi_{2i-1} + \chi_{2i} = 0$ for each i, and because we then have complete character sums we get that

$$\sum_{\chi} \chi(f) = n^{m/2} 1_{f(1)=f(2),\dots,f(m-1)=f(m)} + O_m(n^{m/2-1}).$$

Similarly, for general N and $\mathcal{P}$ we have

$$\sum_{\chi} \chi(f) = n^{m/2} 1_{f \text{ constant on each } P \in \mathcal{P}} + O_m(n^{m/2-1}). \tag{3.4}$$

Now we must sum over N and $\mathcal{P}$, but by again accepting a negligible error we may simply sum over all ways of choosing $m/2$ pairs from $\{1, \dots, n\}$ *with* replacement. Thus

$$\begin{aligned}
&\sum_{|N|=m} \sum_{\mathcal{P}} 1_{f \text{ constant on each } P \in \mathcal{P}} \\
&\quad = \frac{1}{(m/2)! 2^{m/2}} \sum_{i_1, i_2, \dots, i_{m-1}, i_m} 1_{f(i_1)=f(i_2),\dots,f(i_{m-1})=f(i_m)} + O_m(n^{3m/2-1}) \\
&\quad = \frac{1}{(m/2)! 2^{m/2}} \left(\sum_{i,j} 1_{f(i)=f(j)} \right)^{m/2} + O_m(n^{3m/2-1}) \\
&\quad = \frac{1}{(m/2)! 2^{m/2}} \left(\sum_{x \in G} |f^{-1}(x)|^2 \right)^{m/2} + O_m(n^{3m/2-1}).
\end{aligned} \tag{3.5}$$

Thus by combining (3.3), (3.4), and (3.5) we have

$$\sum_{\chi\in\mathfrak{M}_m}\widehat{1_S}(\chi)^3\chi(f)=\frac{1}{(m/2)!}\left(-\frac{1}{n^2}\sum_{x\in G}|f^{-1}(x)|^2\right)^{m/2}\left(\frac{n!}{n^n}\right)^3 + O_m\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right),$$

and this combined with Lemma 3.2 proves Theorem 3.1.

4. Low-entropy minor arcs

In this section we establish a suitable replacement for Theorem 2.2 in the even-order case. Our only tool will be the following recursive formula for $\widehat{1_S}(\chi)$, which was also our only tool for proving Theorem 2.2. Let $\chi\in\widehat{G}^n$ be a character with exactly m nonzero coordinates, say

$$\chi=(\chi_1,\dots,\chi_m,0^{n-m}),$$

where each χ_i is nonzero. Then

$$\begin{aligned}\widehat{1_S}(\chi)&=\frac{(n-m)!}{n^n}\sum_{x_1,\dots,x_{m-1}}{}'\sum_{x_m\neq x_1,\dots,x_{m-1}}\chi_1(x_1)\cdots\chi_m(x_m)\\&=-\frac{(n-m)!}{n^n}\sum_{x_1,\dots,x_{m-1}}{}'\sum_{x_m=x_1,\dots,x_{m-1}}\chi_1(x_1)\cdots\chi_m(x_m)\\&=-\frac{1}{n-m+1}\sum_{i=1}^{m-1}\widehat{1_S}(\chi^i),\end{aligned}\tag{4.1}$$

where here we write χ^i for the character

$$\chi^i=(\chi_1,\dots,\chi_i+\chi_m,\dots,\chi_{m-1},0^{n-m+1}).$$

Note that each χ^i has either $m-1$ or $m-2$ nonzero coordinates (depending on whether $\chi_i=-\chi_m$), so repeated application of (4.1) (and $\widehat{1_S}(0)=n!/n^n$) constitutes a method of computing $\widehat{1_S}(\chi)$. The more relevant thing for us however is the bound implied by (4.1):

$$|\widehat{1_S}(\chi)|\le\frac{1}{n-m+1}\sum_{i=1}^{m-1}|\widehat{1_S}(\chi^i)|.\tag{4.2}$$

An immediate consequence of this bound is the following lemma, which, although weaker than Theorem 2.2 by a factor of roughly $2^{m/2}$, does not rely on the absence of 2-torsion.

Lemma 4.1. *If $\chi\in\widehat{G}^n$ has exactly m nonzero coordinates, where $m\le n/2$, then*

$$|\widehat{1_S}(\chi)|\le\binom{n}{m}^{-1/2}\frac{n!}{n^n}.$$

Proof. The claim is true when $m=0$ (with equality) and when $m=1$ (because $\widehat{1_S}(\chi)=0$), so assume $m>1$. By (4.2),

$$|\widehat{1_S}(\chi)|\le\frac{m-1}{n-m+1}\max_i|\widehat{1_S}(\chi^i)|.$$

Since each χ^i has either $m-1$ or $m-2$ nonzero coordinates, by induction (and $m \leq n/2$) we have

$$\begin{aligned}|\widehat{1_S}(\chi)| &\leq \frac{m-1}{n-m+1}\binom{n}{m-2}^{-1/2}\frac{n!}{n^n} \\ &\leq \left(\frac{m(m-1)}{(n-m)(n-m+1)}\right)^{1/2}\binom{n}{m-2}^{-1/2}\frac{n!}{n^n} \\ &= \binom{n}{m}^{-1/2}\frac{n!}{n^n}. \end{aligned}$$

□

Unfortunately, there are many approximate equality cases of Lemma 4.1. For example if

$$\chi = (\chi_0^m, 0^{n-m}),$$

where $2\chi_0 = 0$, then

$$|\widehat{1_S}(\chi)| = \frac{(m-1)(m-3)\cdots 1}{(n-m+1)(n-m+3)\cdots n}\frac{n!}{n^n} \approx \binom{n}{m}^{-1/2}\frac{n!}{n^n}.$$

This follows from (4.1). Moreover we also expect approximate equality whenever χ approximately has this form, say whenever

$$\chi = (\chi_0^{m-k}, \chi_1, \ldots, \chi_k, 0^{n-m})$$

with $2\chi_0 = 0$ and k much smaller than m. However, our next theorem asserts that we have an exponential improvement to Lemma 4.1 whenever χ does not have essentially this form. This is our replacement for Theorem 2.2.

Theorem 4.2. *Let $\chi \in \widehat{G}^n$ be a character with exactly m nonzero coordinates, where $m \leq n/2$, and such that no nonzero 2-torsion coordinate is repeated more than $(1-\delta)m$ times, where $0 \leq \delta \leq 1/2$. Then*

$$|\widehat{1_S}(\chi)| \leq e^{O(m^{3/2}/n^{1/2}+m^{1/2})}(1-\delta)^{m/2}\binom{n}{m}^{-1/2}\frac{n!}{n^n}.$$

Proof. Fixing δ, let X_m be the set of $\chi \in \widehat{G}^n$ with exactly m nonzero coordinates and such that no nonzero 2-torsion coordinate is repeated more than $(1-\delta)m+3$ times. This set X_m includes all the χ covered by the theorem. (We allow the extra 3 for technical reasons to do with the induction.) Let

$$U_m = \max_{\chi \in X_m} |\widehat{1_S}(\chi)|.$$

We seek a bound for U_m.

Let $\chi \in X_m$. We may assume by permuting coordinates that

$$\chi = (\chi_1, \ldots, \chi_m, 0^{n-m}).$$

We also assume, by permuting coordinates if necesary, the following: if χ has a majority coordinate (i.e., if more than half of the nonzero coordinates $\chi_1, \ldots, \chi_m$ are equal to some particular χ_0) then χ_m is the majority coodinate.

Now suppose we apply (4.2). Consider the resulting characters χ^i:

$$\chi^i = (\chi_1, \ldots, \chi_i + \chi_m, \ldots, \chi_{m-1}, 0^{n-m+1}).$$

We claim that $\chi^i \in X_{m-1} \cup X_{m-2}$. To see this, note first that coordinates other than χ_m are repeated in χ^i at most

$$m/2 + 1 \leq (1-\delta)(m-2) + 3$$

times, so we need only worry about χ_m. If $2\chi_m \neq 0$ then there is nothing to worry about, so assume $2\chi_m = 0$. Suppose χ_m repeats k times in χ. Then depending on whether $\chi_i = \chi_m$, χ^i has either $m-1$ nonzero coordinates and $k-1$ repetitions of χ_m, or $m-2$ nonzero coordinates and $k-2$ repetitions of χ_m. Both of these possibilities are fine: since $k \leq (1-\delta)m + 3$ we have

$$k - 1 \leq (1-\delta)(m-1) + 3,$$
$$k - 2 \leq (1-\delta)(m-2) + 3.$$

Thus indeed either $\chi^i \in X_{m-1}$ or $\chi^i \in X_{m-2}$.

For how many i can we have $\chi^i \in X_{m-2}$? Recall that $\chi^i \in X_{m-2}$ if and only if $\chi_i = -\chi_m$. If more than half of the indices i satisfy $\chi_i = -\chi_m$ then $-\chi_m$ is a majority coordinate. By arrangement, this can only happen if $-\chi_m = \chi_m$, i.e., if χ_m is 2-torsion. By definition of X_m, then, χ_m is repeated at most $(1-\delta)m + 3$ times. Thus there can be at most $(1-\delta)m + 2$ indices i for which $\chi^i \in X_{m-2}$.

Putting this all together with (4.2), we have

$$U_m \leq \frac{m-1}{n-m+1} \max\left(U_{m-1}, \frac{\delta m - 3}{m-1} U_{m-1} + \frac{(1-\delta)m+2}{m-1} U_{m-2}\right) \tag{4.3}$$

for all $m \geq 2$. This recurrence is very similar to the one we analyzed in [EMM15, Section 6]. We use the same method here.

In (4.3), the max covers the possibility that $U_{m-1} > U_{m-2}$, but typically we expect U_{m-1} to be much smaller than U_{m-2}, so in fact only the term involving U_{m-2} should be important. To take advantage of this heuristic, we renormalize. Write

$$\alpha_m = \frac{(1-\delta)m+2}{n-m+1}, \quad \beta_m = \frac{\delta m - 3}{n-m+1}, \quad \gamma_m = \frac{m-1}{n-m+1},$$

and define

$$V_m = (\alpha_1 \alpha_2 \dots \alpha_m)^{-1/2} U_m.$$

Then (4.3) becomes

$$V_m \leq \max\left(\gamma_m \alpha_m^{-1/2} V_{m-1},\ \alpha_m^{1/2} \alpha_{m-1}^{-1/2} V_{m-2} + \beta_m \alpha_m^{-1/2} V_{m-1}\right). \tag{4.4}$$

If we now define

$$v_m = \begin{pmatrix} V_m \\ V_{m-1} \end{pmatrix},$$
$$M_m = \begin{pmatrix} \gamma_m \alpha_m^{-1/2} & 0 \\ 1 & 0 \end{pmatrix},$$
$$N_m = \begin{pmatrix} \beta_m \alpha_m^{-1/2} & \alpha_m^{1/2} \alpha_{m-1}^{-1/2} \\ 1 & 0 \end{pmatrix},$$

then we can write (4.4) equivalently as

$$\|v_m\|_2 \leq \max\left(\|M_m v_{m-1}\|_2, \|N_m v_{m-1}\|_2\right).$$

To finish we bound the operator norms

$$\|M_m\|_{2\to 2}, \quad \|N_m\|_{2\to 2},$$

and then use the bound

$$\begin{aligned}|V_m| &\leq \|v_m\|_2 \\ &\leq \prod_{r=2}^{m} \max\left(\|M_m\|_{2\to 2}, \|N_m\|_{2\to 2}\right) \|v_1\|_2 \\ &= \prod_{r=2}^{m} \max\left(\|M_m\|_{2\to 2}, \|N_m\|_{2\to 2}\right) \frac{n!}{n^n}. \end{aligned} \tag{4.5}$$

Since $\alpha_m \alpha_{m-1}^{-1} = 1 + O(1/m)$ and $\beta_m^2 \alpha_m^{-1} = O(m/n)$, we have

$$\operatorname{tr}\left(N_m^T N_m\right) = 2 + O(m/n + 1/m)$$

and

$$\det\left(N_m^T N_m\right) = 1 + O(1/m).$$

Thus (because $\|N_m\|_{2\to 2}^2$ is equal to the larger eigenvalue of $N_m^T N_m$) we have

$$\|N_m\|_{2\to 2} = 1 + O(m^{1/2}/n^{1/2} + 1/m^{1/2}).$$

Similarly, because $\gamma_m^2 \alpha_m^{-1} = O(m/n)$, we have

$$\|M_m\|_{2\to 2} = 1 + O(m/n).$$

Thus from (4.5) we have

$$\begin{aligned}|V_m| &\leq \prod_{r=2}^{m} \left(1 + O(r^{1/2}/n^{1/2} + 1/r^{1/2})\right) \frac{n!}{n^n} \\ &= e^{O(m^{3/2}/n^{1/2} + m^{1/2})} \frac{n!}{n^n}. \end{aligned}$$

Since

$$\alpha_1 \cdots \alpha_m = \prod_{r=1}^{m} \left((1 + O(1/r)) \frac{(1-\delta)r}{n-r+1}\right) = e^{O(\log m)} (1-\delta)^m \binom{n}{m}^{-1},$$

the theorem follows. □

Corollary 4.3. *Assume $0 \leq \delta \leq 1/2$. For $m \leq n/2$ we have*

$$\begin{aligned}\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \leq e^{O(m^{3/2}/n^{1/2} + m^{1/2})} (1-\delta)^{m/2} \left(\frac{n!}{n^n}\right)^3 \\ + 2^m n^{\delta m + 1} \binom{n}{m}^{-1/2} \left(\frac{n!}{n^n}\right)^3. \end{aligned}$$

Thus, for some constants $\eta < 1$ and M_0, for all $M \geq M_0$ we have

$$\sum_{m=M}^{n/(\log n)^2} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \leq O\left(\eta^M \left(\frac{n!}{n^n}\right)^3\right) + O\left(e^{-cn^{1/2}} \left(\frac{n!}{n^n}\right)^3\right).$$

Proof. Again denote by X_m the set of characters to which Theorem 4.2 applies. By combining Theorem 4.2 with Theorem 2.3 we have

$$\begin{aligned}\sum_{\chi \in X_m} |\widehat{1_S}(\chi)|^3 &\leq \max_{\chi \in X_m} |\widehat{1_S}(\chi)| \sum_{\chi \in X_m} |\widehat{1_S}(\chi)|^2 \\ &\leq e^{O(m^{3/2}/n^{1/2} + m^{1/2})} (1-\delta)^{m/2} \left(\frac{n!}{n^n}\right)^3. \end{aligned}$$

This is the first term in the claim. On the other hand the number of m-sparse χ not in X_m is at most

$$\binom{n}{m}2^m n^{\delta m+1}$$

(you choose m coordinates to be nonzero, some subset of at most δm of these to be different from χ_0, and finally the values of the coordinates). Thus by Lemma 4.1 we have

$$\sum_{\substack{m\text{-sparse } \chi \\ \chi\notin X_m}} |\widehat{1_S}(\chi)|^3 \leq 2^m n^{\delta m+1}\binom{n}{m}^{-1/2}\left(\frac{n!}{n^n}\right)^3.$$

To prove the second part of the corollary, consider the second term first. Using the simple bound $\binom{n}{m} \geq (n/m)^m$, we have

$$\begin{aligned}\log\left(2^m n^{\delta m+1}\binom{n}{m}^{-1/2}\right) &\leq \log\left(2^m n^{\delta m+1}(n/m)^{-m/2}\right)\\ &= O(m+\log n)+\delta m\log n-(m/2)\log(n/m).\end{aligned}$$

As long as n/m and m are both larger than a sufficiently large constant (depending on the constant implicit in the O above), then $m+\log n$ is negligible compared to $(m/2)\log(n/m)$, so if we put

$$\delta = c\log(n/m)/\log n$$

for a sufficiently small constant c then we have

$$\log\left(2^m n^{\delta m+1}\binom{n}{m}^{-1/2}\right) \leq -(m/4)\log(n/m).$$

Now consider the first term. we have

$$\begin{aligned}\log\left(e^{O(m^{3/2}/n^{1/2}+m^{1/2})}(1-\delta)^{m/2}\right) &\leq O(m^{3/2}/n^{1/2}+m^{1/2})-\delta m/2\\ &= \left(-(c/2)\frac{\log(n/m)}{\log n}+O((m/n)^{1/2}+m^{-1/2})\right)m.\end{aligned}$$

As long as $1 \leq m = o(n(\log\log n)^2/(\log n)^2)$ it is easy to see that

$$(m/n)^{1/2} = o\left(\frac{\log(n/m)}{\log n}\right):$$

indeed, this is equivalent to

$$\log n = o((n/m)^{1/2}\log(n/m)).$$

It is also easy to see that

$$m^{-1/2} = o\left(\frac{\log(n/m)}{\log n}\right)$$

as long as $m\to\infty$ (consider the cases $m\leq n^{1/2}$ and $m\geq n^{1/2}$ separately). Thus we have

$$\log\left(e^{O(m^{3/2}/n^{1/2}+m^{1/2})}(1-\delta)^{m/2}\right) \leq -(c/4)\frac{\log(n/m)}{\log n}m$$

as long as $m\geq M_0$ for some constant $M_0 = M_0(c)$.

Putting these two bounds together, as long as $M_0 \leq m \leq n/(\log n)^2$ we have

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \leq \exp\left(-c'm\frac{\log(n/m)}{\log n}\right)\left(\frac{n!}{n^n}\right)^3.$$

In particular if $m \geq n^{0.6}$ we have

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \leq \exp\left(-c'n^{0.6}/\log n\right)\left(\frac{n!}{n^n}\right)^3.$$

In this range we can afford to simply take the maximum over m and accept an additional factor of n. On the other hand if $m \leq n^{0.6}$ we have

$$\sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3 \leq \exp\left(-0.6c'm\right)\left(\frac{n!}{n^n}\right)^3.$$

The sum of this expression over $m \geq M$ is a geometric series and thus dominated by the $m = M$ term. This proves the corollary. □

5. High-entropy minor arcs

In this section we extend the proof of (2.2) to cover a wider entropy range. This mainly requires more careful counting. The methods of this section are effective for characters with entropy as low as $n^{-1/2+\varepsilon}$, but for simplicity we only consider entropy as low as $(\log n)^{-100}$.

We first need a lemma relating our definition of entropy to the usual one.

Lemma 5.1. *Let $H = H(\chi) = \frac{1}{n}\log\binom{n}{a_1,\dots,a_n}$. If $H \geq (\log n)^{-100}$ then*

$$H = \left(1 + O\left(\frac{\log\log n}{\log n}\right)\right)\sum_{i=1}^{n}\frac{a_i}{n}\log\frac{n}{a_i}.$$

Proof. We use Stirling's formula in the form

$$\log m! = m\log m - m + O(\log(m+2)) \qquad (m \geq 0).$$

This implies

$$\frac{1}{n}\log\binom{n}{a_1,\dots,a_n} = \sum_{a_i>0}\left(\frac{a_i}{n}\log\frac{n}{a_i} + O\left(\frac{\log(a_i+2)}{n}\right)\right) + O\left(\frac{\log n}{n}\right). \quad (5.1)$$

Suppose k of the a_i are nonzero. Note first of all that (5.1) implies a bound on k: we have

$$\begin{aligned} H &= \frac{1}{n}\log\binom{n}{a_1,\dots,a_n} \\ &\geq \frac{1}{n}\log\left(\frac{n!}{(n-k+1)!1^{k-1}}\right) \\ &= \frac{n-k+1}{n}\log\frac{n}{n-k+1} + \frac{k-1}{n}\log n + O\left(\frac{k}{n} + \frac{\log n}{n}\right) \\ &= \frac{k}{n}\log n + O\left(\frac{k}{n} + \frac{\log n}{n}\right). \end{aligned}$$

This implies $k \leq O(Hn/\log n)$. Thus by concavity of log we have

$$\begin{aligned}\frac{1}{n}\sum_{a_i>0}\log(a_i+2) &\leq \frac{k}{n}\log\left(\frac{1}{k}\sum_{a_i>0}(a_i+2)\right)\\ &= \frac{k}{n}\log\left(\frac{n}{k}+2\right)\\ &\leq O\left(\frac{H}{\log n}\left(\log\log n+\log H^{-1}\right)\right)\\ &= O\left(\frac{H\log\log n}{\log n}\right).\end{aligned}$$

In the last line we used the assumption $H \geq (\log n)^{-100}$. Inserting this into (5.1) gives

$$H = \sum_{i=1}^{n}\frac{a_i}{n}\log\frac{n}{a_i} + O\left(\frac{H\log\log n}{\log n}\right),$$

which is equivalent to the claim. □

Lemma 5.2. *Suppose $H \geq (\log n)^{-100}$. Then the number of characters $\chi \in \widehat{G}^n$ of entropy at most H is bounded by*

$$\exp\left(Hn + O\left(\frac{Hn\log\log n}{\log n}\right)\right).$$

The important part of the conclusion is that the bound has the form $e^{Hn+o(Hn)}$.

Proof. If χ has entropy H then the orbit of χ under permutation of coordinates has size e^{Hn}, so it suffices to show that the number of orbits is at most

$$\exp\left(O\left(\frac{Hn\log\log n}{\log n}\right)\right).$$

Choosing an orbit is equivalent to fixing the multiplicities a_i of each coordinate, so we are in the business of counting solutions to

$$\begin{cases} a_1+\cdots+a_n & = n,\\ \frac{1}{n}\log\binom{n}{a_1,\ldots,a_n} & \leq H.\end{cases} \tag{5.2}$$

Let t be the sum of the a_i for which $a_i \leq n^{1/2}$, and note that at most $n^{1/2}$ of the a_i are not included in this set. Let

$$H' = \sum_{i=1}^{n}\frac{a_i}{n}\log\frac{n}{a_i}.$$

Then

$$H' \geq \frac{t}{n}\log n^{1/2} = \frac{t}{2n}\log n.$$

On the other hand Lemma 5.1 implies $H' \leq O(H)$. Thus

$$t \leq O\left(\frac{Hn}{\log n}\right).$$

Now we can select a solution to (5.2) by choosing up to $n^{1/2}$ indices i at which we will have $a_i > n^{1/2}$, choosing the values of these a_i, and then choosing the values

of all other a_i in such a way that their sum does not exceed $O(Hn/\log n)$. Thus the number of solutions to (5.2) is at most

$$n^{O(n^{1/2})}\binom{n+O(Hn/\log n)}{O(Hn/\log n)} = n^{O(n^{1/2})}\exp\left(O\left(\frac{Hn}{\log n}\log\frac{\log n}{H}\right)\right)$$
$$= \exp\left(O\left(\frac{\log\log n}{\log n}+\frac{\log H^{-1}}{\log n}\right)Hn\right)$$
$$= \exp\left(O\left(\frac{Hn\log\log n}{\log n}\right)\right). \qquad \square$$

Theorem 5.3. *We have the following bound over high-entropy characters:*

$$\sum_{H(\chi)\geq(\log n)^{-100}} |\widehat{1_S}(\chi)|^3 \leq e^{-cn(\log n)^{-100}}\left(\frac{n!}{n^n}\right)^3.$$

Proof. Suppose $\chi \in \widehat{G}^n$ is a character of entropy at least $(\log n)^{-100}$. By Theorem 2.4 we have

$$|\widehat{1_S}(\chi)| \leq \binom{n+k-1}{k-1}^{1/2} e^{-Hn/2}\frac{n!}{n^n},$$

where k is the number of distinct coordinates of χ. We saw in the proof of Lemma 5.1 that $k \leq O(Hn/\log n)$. Thus

$$\binom{n+k-1}{k-1} \leq \exp\left(O\left(\frac{Hn\log\log n}{\log n}\right)\right) = \exp(o(Hn)).$$

We therefore have

$$\sum_{H(\chi)\geq(\log n)^{-100}} |\widehat{1_S}(\chi)|^3 \leq \sum_{H(\chi)\geq(\log n)^{-100}} e^{-3H(\chi)n/2+o(H(\chi)n)}\left(\frac{n!}{n^n}\right)^3.$$

Thus by applying Lemma 5.2 and a dyadic decomposition on H we have

$$\sum_{H(\chi)\geq(\log n)^{-100}} |\widehat{1_S}(\chi)|^3 \leq \sum_{\substack{(\log n)^{-100}\leq H\leq\log n\\ H \text{ dyadic}}} e^{-Hn/2+o(Hn)}\left(\frac{n!}{n^n}\right)^3$$
$$\leq O(\log\log n)e^{-cn(\log n)^{-100}}\left(\frac{n!}{n^n}\right)^3. \qquad \square$$

6. Proofs of the main theorems

We have now assembled all the tools we need to prove our main theorems. We closely follow the outline from Section 2.

Proof of Theorem 1.3. As explained in Section 2, the number of solutions to $\pi_1 + \pi_2 + \pi_3 = f$ is

$$n^{2n}\sum_{\chi\in\widehat{G}^n} \widehat{1_S}(\chi)^3\chi(f),$$

and we have the approximation

$$\sum_{\chi\in\widehat{G}^n}\widehat{1_S}(\chi)^3\chi(f) = n\sum_{m=0}^{M}\sum_{m\text{-sparse }\chi}\widehat{1_S}(\chi)^3\chi(f) + O\left(n\sum_{m=M+1}^{2\varepsilon n}\sum_{m\text{-sparse }\chi}|\widehat{1_S}(\chi)|^3\right) + O\left(\sum_{H(\chi)>\varepsilon}|\widehat{1_S}(\chi)|^3\right), \tag{6.1}$$

provided that $\sum_{i=1}^n f(i) = \sum G$. Put $\varepsilon = (\log n)^{-100}$. As long as $M \geq M_0$, Corollary 4.3 implies that the second term of (6.1) has size

$$O\left(\eta^M n\left(\frac{n!}{n^n}\right)^3\right) + O\left(e^{-cn^{1/2}}\left(\frac{n!}{n^n}\right)^3\right),$$

while Theorem 5.3 implies that the last term has size

$$O\left(e^{-cn(\log n)^{-100}}\left(\frac{n!}{n^n}\right)^3\right).$$

Meanwhile, by Theorem 3.1 we have

$$\sum_{m=0}^{M}\sum_{m\text{-sparse }\chi}\widehat{1_S}(\chi)^3\chi(f) = \sum_{\substack{m\leq M\\ m\text{ even}}}\frac{1}{(m/2)!}\left(-\frac{1}{2n^2}\sum_{x\in G}|f^{-1}(x)|^2\right)^{m/2}\left(\frac{n!}{n^n}\right)^3 + O_M\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)^3\right).$$

Thus if $M\to\infty$ sufficiently slowly as a function of n then we have

$$\sum_{m=0}^{M}\sum_{m\text{-sparse }\chi}\widehat{1_S}(\chi)^3\chi(f) = \left(\exp\left(-\frac{1}{2n^2}\sum_{x\in G}|f^{-1}(x)|^2\right)+o(1)\right)\left(\frac{n!}{n^n}\right)^3.$$

This proves the theorem. □

The proof of Theorem 1.4 is similar, but a little easier because we don't need the careful major arc calculations.

Proof of Theorem 1.4. The number of solutions to

$$\pi_1+\cdots+\pi_d = f$$

is

$$n^{(d-1)n}1_S^{*d}(f) = n^{(d-1)n}\sum_{\chi\in\widehat{G}^n}\widehat{1_S}(\chi)^d\chi(f).$$

Let $\widehat{G}_d\subset\widehat{G}^n$ be the set of characters of the form $\chi=(\chi_0^n)$. Then

$$\sum_{\chi\in\widehat{G}_d}\widehat{1_S}(\chi)^d\chi(f) = \sum_{\chi_0\in\widehat{G}}\widehat{1_S}(\chi_0^n)^d\chi_0(\textstyle\sum f) = \displaystyle\sum_{\chi_0\in\widehat{G}}\chi_0(d\textstyle\sum G+\sum f)\left(\frac{n!}{n^n}\right)^d.$$

Assuming $\sum_{i=1}^n f(i) = d\sum G$, this is $n(n!/n^n)^d$.

On the other hand, for $d \geq 4$ we have

$$\sum_{\chi \notin \widehat{G}_d} |\widehat{1_S}(\chi)|^d \leq \max_{\chi \notin \widehat{G}_d} |\widehat{1_S}(\chi)|^{d-3} \sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^3.$$

It follows from Lemma 4.1 and shift-invariance of $|\widehat{1_S}(\chi)|$ (see (2.1)) that

$$\max_{\chi \notin \widehat{G}_d} |\widehat{1_S}(\chi)| \leq \binom{n}{2}^{-1/2} \frac{n!}{n^n} = O\left(\frac{1}{n}\left(\frac{n!}{n^n}\right)\right).$$

Meanwhile we can decompose the sum

$$\sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^3.$$

as

$$O\left(n \sum_{m=0}^{M} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3\right) + O\left(n \sum_{m=M+1}^{2\varepsilon n} \sum_{m\text{-sparse } \chi} |\widehat{1_S}(\chi)|^3\right) + O\left(\sum_{H(\chi)>\varepsilon} |\widehat{1_S}(\chi)|^3\right).$$

The second two terms are bound exactly as in the previous proof (i.e., by the second part of Corollary 4.3 and Theorem 2.2), while the first can be bound using the first part of Corollary 4.3. The result is that[1]

$$\sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^3 \leq O\left(n\left(\frac{n!}{n^n}\right)^3\right).$$

Thus

$$\max_{\chi \notin \widehat{G}_d} |\widehat{1_S}(\chi)|^{d-3} \sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^3 \leq O_d\left(n^{4-d}\left(\frac{n!}{n^n}\right)\right),$$

and the theorem is proved. □

Finally we prove Theorem 1.5.

Proof of Theorem 1.5. We have

$$\left\| \left(\frac{n^m(n-m)!}{n!}\right)^2 1_{S_m} * 1_{S_m} - 1 \right\|_2^2 = \left(\frac{n^m(n-m)!}{n!}\right)^4 \sum_{\substack{\chi \in \widehat{G}^m \\ \chi \neq 0}} |\widehat{1_{S_m}}(\chi)|^4. \quad (6.2)$$

The Fourier transform of 1_{S_m} is related to that of 1_S by

$$\widehat{1_{S_m}}(\chi_1, \dots, \chi_m) = \frac{n^{n-m}}{(n-m)!} \widehat{1_S}(\chi_1, \dots, \chi_m, 0^{n-m}).$$

[1] In fact a finer analysis using major arc calculations proves that

$$\sum_{\chi \in \widehat{G}^n} |\widehat{1_S}(\chi)|^3 = (e^{1/2} + o(1))n\left(\frac{n!}{n^n}\right)^3,$$

but we don't need this.

Thus we can expand (6.2) as

$$\left(\frac{n^m(n-m)!}{n!}\right)^4 \sum_{\substack{\chi\in\widehat{G}^m\\ \chi\neq 0}} |\widehat{1_{S_m}}(\chi)|^4 = \left(\frac{n^n}{n!}\right)^4 \sum_{\substack{\chi\in\widehat{G}^m\\ \chi\neq 0}} |\widehat{1_S}(\chi_1,\dots,\chi_m,0^{n-m})|^4.$$

Now it is easy to see (say by combining Lemma 4.1 with some basic major arc calculations) that we have

$$|\widehat{1_S}(\chi)| \leq O\left(\frac{1}{n^3}\left(\frac{n!}{n^n}\right)\right)$$

unless χ is, up to translation, at most 4-sparse. Let X be the set of these χ. Then

$$\sum_{\chi\notin X} |\widehat{1_S}(\chi)|^4 \leq O\left(\frac{1}{n^4}\left(\frac{n!}{n^n}\right)\right)\sum_{\chi} |\widehat{1_S}(\chi)|^3 \leq O\left(\frac{1}{n^3}\left(\frac{n!}{n^n}\right)^4\right),$$

since as we saw in the previous proof we have $\sum_\chi |\widehat{1_S}(\chi)|^3 \leq O(n(n!/n^n)^3)$. Finally, it is straightforward to calculate that, as long as $m \leq n-1$, we have

$$\left(\frac{n^n}{n!}\right)^4 \sum_{\substack{\chi\in\widehat{G}^m\\ (\chi,0^{n-m})\in X,\chi\neq 0}} |\widehat{1_S}(\chi_1,\dots,\chi_m,0^{n-m})|^4 = O(m^2/n^3),$$

the dominant contribution coming from 2-sparse characters, and this proves the claim.

The very last part of the theorem, about total variation distance, holds because the L^2 distance bounds the L^1 distance, and the L^1 distance is exactly twice the total variation distance. □

7. Transversals in Latin hypercubes

A *Latin square of order n* is an $n\times n$ grid populated with the symbols $\{1,\dots,n\}$ in such a way that no symbol occurs more than once in any given row or column. A *transversal* in a Latin square is a selection of n entries not repeating any row, column, or symbol. See Figure 1 for an example of a transversal in a particular Latin square of order 7. See Wanless [Wan11] for an extensive survey about transversals in Latin squares.

(1)	2	3	4	5	6	7
2	6	1	(7)	3	5	4
6	5	7	2	(4)	1	3
5	4	(6)	3	2	7	1
4	7	2	6	1	3	(5)
3	1	4	5	7	(2)	6
7	(3)	5	1	6	4	2

Figure 1. A Latin square and a transversal

More generally, a *Latin (hyper)cube of dimension d and order n* is a d-dimensional grid of side length n filled out with $\{1,\dots,n\}$ in such a way that no symbol occurs

more than once in the same (axis-aligned) line. A *transversal* in a d-dimensional Latin cube is a selection of n entries not repeating any hyperplane or symbol.

Now let G be an abelian group of order n. We claim that counting solutions to $\pi_1 + \cdots + \pi_d = \pi_{d+1}$ in G is equivalent to counting transversals in a certain d-dimensional Latin cube. Let $\pi : \{1, \dots, n\} \to G$ be a bijection. The *Latin cube* $L^d(G, \pi)$ *induced by* G *and* π is the d-dimensional grid of side length n whose $(i_1, \dots, i_d)$-entry is

$$x = \pi^{-1}\left(\pi(i_1) + \cdots + \pi(i_d)\right). \tag{7.1}$$

Clearly this defines a Latin cube: if we alter any one of the i_j while keeping the others fixed then by the group property we must alter x. One can think of $L^2(G, \pi)$ as the multiplication (or addition) table of G, and likewise $L^d(G, \pi)$ as the graph of the d-wise iterated group operation.

Lemma 7.1. *Let $S \subset G^n$ be the set of bijections $\pi : \{1, \dots, n\} \to G$. The transversals of $L^d(G, \pi)$ are in bijection with the solutions to*

$$\pi_1 + \cdots + \pi_d = \pi \qquad (\pi_1, \dots, \pi_d \in S). \tag{7.2}$$

In particular $T(L^d(G, \pi))n!$ is the number of solutions to

$$\pi_1 + \cdots + \pi_d = \pi_{d+1} \qquad (\pi_1, \dots, \pi_{d+1} \in S). \tag{7.3}$$

Proof. Fix a transversal in $L^d(G, \pi)$. For each symbol $x \in \{1, \dots, n\}$ let

$$(\pi^{-1}\pi_1(x), \dots, \pi^{-1}\pi_d(x)) \tag{7.4}$$

be the coordinates of the symbol x in the transversal. Since each symbol occurs exactly once in a transversal, $\pi_1, \dots, \pi_d$ are well defined functions $\{1, \dots, n\} \to G$. The hyperplane condition for transversals implies that $\pi_1, \dots, \pi_d$ are injections, hence bijections. Finally, the definition (7.1) of $L^d(G, \pi)$ implies (7.2).

Conversely, given a solution $(\pi_1, \dots, \pi_d)$ to (7.2) with $\pi_1, \dots, \pi_d \in S$, consider the collection of entries of $L^d(G, \pi)$ indexed by (7.4). The condition (7.2) ensures that each symbol is represented exactly once (since x is in position (7.4)), and the condition that $\pi_1, \dots, \pi_d \in S$ ensures that no two entries occur in the same hyperplane.

For the last part of the lemma, just note that

$$\pi_1 + \cdots + \pi_d = \pi$$

if and only if

$$\pi_1\pi^{-1}\pi_{d+1} + \cdots + \pi_d\pi^{-1}\pi_{d+1} = \pi_{d+1}.$$

Thus every solution to (7.2) corresponds to $n!$ solutions to (7.3). □

Thus by applying Theorems 1.2 and 1.4 we have the following.

Theorem 7.2. *Let G be an abelian group of order n. Then the number of transversals in $L^d(G, \pi)$ is*

$$\begin{cases} (e^{-1/2} + o(1))n!^2/n^{n-1} & \text{if } d = 2 \text{ and } \sum G = 0, \\ (1 + O_d(n^{2-d}))n!^d/n^{n-1} & \text{if } d \geq 3 \text{ and } (d+1)\sum G = 0, \\ 0 & \text{otherwise.} \end{cases}$$

For context in the Latin squares literature, let $T(d,n)$ be the maximum number of transversals in a Latin cube of dimension d and order n. Taranenko [Tar15, Theorem 6.1] proved that

$$T(d,n) \leq \left((1+o(1))\frac{n^{d-1}}{e^d}\right)^n.$$

A matching lower bound was proved by Glebov and Luria [GL15, Theorems 1.3 and 4.1] using a random construction: they proved that

$$T(d,n) \geq \left((1-o(1))\frac{n^{d-1}}{e^d}\right)^n$$

for all n when $d=2$, and for infinitely many n for each $d \geq 3$. Theorem 7.2 implies a more precise lower bound for $T(d,n)$, with an explicit Latin cube, whenever d is odd or $n \not\equiv 2 \pmod 4$. It is an amusing challenge to cover the remaining cases, but not one we entertain here.

8. PRP-to-PRF conversion

Two fundamental building blocks in cryptography are so-called *pseudorandom functions* (PRFs) and *pseudorandom permutations* (PRPs). Roughly speaking, a pseudorandom function is a family of functions $f_s : \{0,1\}^d \to \{0,1\}^d$ such that

- for any s and $x \in \{0,1\}^d$ there is an efficient algorithm for computing $f_s(x)$,
- if s is chosen uniformly at random then it is not computationally practical to distinguish f_s from a truly random function $\{0,1\}^d \to \{0,1\}^d$, with any statistical significance.

Similarly, a pseudorandom permutation is a family of permutation $f_s : \{0,1\}^d \to \{0,1\}^d$ such that

- for any s and $x \in \{0,1\}^d$ there is an efficient algorithm for computing $f_s(x)$,
- if s is chosen uniformly at random then it is not computationally practical to distinguish f_s from a truly random permutation $\{0,1\}^d \to \{0,1\}^d$, with any statistical significance.

Usually the domain of s is taken to be $\{0,1\}^k$ for some k. In both cases we will write $F : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^d$ for the function defined by $F(s,x) = f_s(x)$, and we will refer to F itself as the PRF/PRP. Both concepts are treated as cryptographic primitives, from which more complex cryptographic constructions are built, but they have somewhat different particular use cases. For more detail and background, see Goldreich [Gol01, Chapter 3].

The security of a PRF or a PRP is measured by a hypothetical attacker's *advantage*, which is defined as follows. Suppose A is some probabilistic algorithm which, given a function $f : \{0,1\}^d \to \{0,1\}^d$, outputs either $A(f) = 0$ or $A(f) = 1$. Then the *advantage* of A as a PRF distinguisher is

$$\mathrm{Adv}_F^{\mathrm{PRF}}(A) = |\mathbf{P}(A(f_s) = 1) - \mathbf{P}(A(f) = 1)|,$$

where s is drawn uniformly from $\{0,1\}^k$ and f is an a genuinely uniform random function. The advantage to an attacker with access to at most m queries is then defined as

$$\mathrm{Adv}_F^{\mathrm{PRF}}(m) = \max \mathrm{Adv}_F^{\mathrm{PRF}}(A),$$

where the maximum is taken over all probabilitic algorithms A which query at most m values $f(x)$, where the places x being queried may depend on previous values

(an "adaptive chosen plaintext attack"). Usually some constraint is also placed on the computational power of the attacker, but we can mostly focus on the number of queries. Advantage for PRP attackers is defined similarly.

Given a good PRF generator, a corresponding PRP generator can be constructed using the so-called Feistel cipher. The security of this construction has been extensively studied, starting with the seminal result of Luby and Rackoff [LR88], which states that 3 rounds of the Feistel cipher are sufficient to guarantee security. More speficially, Luby and Rackoff proved that if a 3-round Feistel cipher is combined with an ideal PRF generator then

$$\mathrm{Adv}_F^{\mathrm{PRP}}(m) \leq m^2/2^{d/2}.$$

Moreover we get stronger bounds if more rounds are used: see Patarin [Pat03].

Conversion in the other direction is less well studied, partly because it often suffices just to consider a PRP itself as a PRF. Specifically, if we attempt to use a PRP generator F itself as a PRF, then we have

$$\mathrm{Adv}_F^{\mathrm{PRF}}(m) \leq \mathrm{Adv}_F^{\mathrm{PRP}}(m) + m^2/2^d.$$

However, such a construction is vulnerable to the so-called *birthday attack*: one expects to have to query only about $2^{d/2}$ times before seeing a collision, so roughly $2^{d/2}$ queries should suffice to distinguish a PRP from a PRF. This attack shows that the above bound is close to sharp. We could overcome this problem by, say, doubling the size of d, but this comes at a cost in efficiency.

An alternative construction is to take two independent PRPs and use their bitwise xor as our PRF. In other words we take two independent pseudorandom permutations $\tilde{\pi}_1, \tilde{\pi}_2$ of $\mathbf{F}_2^d$ and we propose $\tilde{\pi}_1 + \tilde{\pi}_2$ as a pseudorandom function. This was first analyzed by Bellare and Impagliazzo [BI99] and Lucks [Luc00], and later by Patarin [Pat13] and others. An optimal security bound for this construction follows from Theorem 1.5.

Theorem 8.1. *Suppose we use the* xor *construction to make a PRF generator F from a PRP generator E. Then as long as $m < 2^d$ we have*

$$\mathrm{Adv}_F^{PRF}(m) \leq 2\mathrm{Adv}_E^{PRP}(m) + O(m/2^{3d/2}).$$

Proof. The following diagram may or may not be helpful:

$$\tilde{\pi}_1 + \tilde{\pi}_2 \xleftarrow{\mathrm{Adv}_E^{\mathrm{PRP}}(m)} \pi_1 + \tilde{\pi}_2 \xleftarrow{\mathrm{Adv}_E^{\mathrm{PRP}}(m)} \pi_1 + \pi_2 \xleftarrow{O(m/2^{3d/2})} f.$$

We are using $\tilde{\pi}_1 + \tilde{\pi}_2$ as our PRF, where $\tilde{\pi}_1$ and $\tilde{\pi}_2$ are independent draws from our PRP generator. The attacker cannot distinguish this from $\pi_1 + \pi_2$, where π_1 and π_2 are truly random permutations $\mathbf{F}_2^d \to \mathbf{F}_2^d$, with advantage greater than $2\mathrm{Adv}_E^{\mathrm{PRP}}(m)$. Since advantage is bounded by total variation distance, Theorem 1.5 implies that the attacker cannot distinguish $\pi_1 + \pi_2$ from a truly random function f with advantage greater than $O(m/2^{3d/2})$. Thus the claim follows from the triangle inequality for advantage. □

For a better introduction to PRP-to-PRF convesion, see Bellare, Krovetz, and Rogaway [BKR98].

Acknowledgements

I am grateful to Freddie Manners and Rudi Mrazović for numerous technical discussions, to Samuel Neves for conversations about cryptography, and to Ian Wanless for conversations about Latin squares.

References

[BI99] M. Bellare and R. Impagliazzo. A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. 1999. http://eprint.iacr.org/1999/024.

[BKR98] M. Bellare, T. Krovetz, and P. Rogaway. Luby-rackoff backwards: Increasing security by making block ciphers non-invertible. In *Advances in Cryptology — EUROCRYPT'98: International Conference on the Theory and Application of Cryptographic Techniques Espoo, Finland, May 31 – June 4, 1998 Proceedings*, pages 266–280. Springer, Berlin, Heidelberg, 1998. http://dx.doi.org/10.1007/BFb0054132.

[Con] K. Conrad. Characters of finite abelian groups. http://www.math.uconn.edu/~kconrad/blurbs/grouptheory/charthy.pdf.

[EMM15] S. Eberhard, F. Manners, and R. Mrazović. Additive triples of bijections, or the toroidal semiqueens problem. 2015. http://arxiv.org/abs/1510.05987.

[GL15] R. Glebov and Z. Luria. On the maximum number of latin transversals. 2015. http://arxiv.org/abs/1506.00983.

[Gol01] O. Goldreich. *Foundations of cryptography*. Cambridge University Press, Cambridge, 2001. Basic tools.

[LR88] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.*, 17(2):373–386, 1988. Special issue on cryptography.

[Luc00] S. Lucks. The sum of PRPs is a secure PRF. In *Advances in cryptology—EUROCRYPT 2000 (Bruges)*, volume 1807 of *Lecture Notes in Comput. Sci.*, pages 470–484. Springer, Berlin, 2000.

[Pat03] J. Patarin. Luby-Rackoff: 7 rounds are enough for $2^{n(1-\varepsilon)}$ security. In *Advances in cryptology—CRYPTO 2003*, volume 2729 of *Lecture Notes in Comput. Sci.*, pages 513–529. Springer, Berlin, 2003.

[Pat13] J. Patarin. Security in $O(2^n)$ for the xor of two random permutations – proof with the standard H technique –. 2013. https://eprint.iacr.org/2013/368.

[Tar15] A. A. Taranenko. Multidimensional permanents and an upper bound on the number of transversals in Latin squares. *J. Combin. Des.*, 23(7):305–320, 2015.

[Wan11] I. M. Wanless. Transversals in Latin squares: a survey. In *Surveys in combinatorics 2011*, volume 392 of *London Math. Soc. Lecture Note Ser.*, pages 403–437. Cambridge Univ. Press, Cambridge, 2011.

Sean Eberhard, London, UK
E-mail address: eberhard.math@gmail.com